DYNAMIC COMPRESSION OF EXPERTISE

Artificial intelligence, competence-output decoupling, and the devaluation of cognitive work

Xema Gatnau

Abstract

Generative artificial intelligence enables individuals with very different abilities, education, and experience to produce outputs of increasingly similar quality. Recent literature has documented particularly large improvements among novice workers and lower-performing students, compression in the distributions of productivity and grades, and a decline in the ability of academic outputs to signal their authors' underlying competences. These phenomena, however, are usually studied statically, by observing the effect of a given technology on an already-formed population.

This article proposes a theory of the dynamic compression of expertise. The model integrates two simultaneous temporal processes. On the one hand, artificial intelligence improves and progressively assumes not only task execution, but also problem formulation, contextualization, verification, and correction—functions that initially protected the expert's advantage. On the other hand, new cohorts are educated using increasingly capable systems, which may reduce the autonomous practice required to develop certain internal competences. The interaction between these two processes produces a growing decoupling between the observable quality of an output and the ability that the individual can demonstrate without assistance.

The theory predicts successive compression between students of different levels, between specialists and generalists, and ultimately between experts and non-experts across a growing share of cognitive tasks. It does not claim that all differences in knowledge disappear, but that they may cease to translate into differences that are visible, scarce, and economically remunerable. Expertise would retain special value in novel, ambiguous, exceptional, or high-responsibility problems, while losing relative value in ordinary cognitive production. The article presents a formal model, testable hypotheses, boundary conditions, and an empirical agenda designed to distinguish productive improvement, internal learning, borrowed performance, and signaling value.

Keywords: generative artificial intelligence; expertise; human capital; higher education; signaling; deskilling; cognitive offloading; productivity; generalists; automation.

JEL classification: I23, J24, O33, D83.

1. Introduction

The widespread introduction of generative artificial intelligence into education and cognitive work is altering a relationship that could long be taken approximately for granted: the relationship between what a person knows and the quality of the output that person is capable of delivering.

An excellent academic paper, a technically sound report, or a well-argued professional response functioned as imperfect but reasonably informative indicators of the competence of the person who produced them. The person could consult books, databases, and other

tools, but the selection, structuring, integration, drafting, and verification of the result depended to a large extent on capacities that had to be internally possessed.

Artificial intelligence changes this relationship. An individual can produce an output that incorporates knowledge, operations, and forms of reasoning that the individual could not reproduce autonomously. The output is not necessarily false or fraudulent. It may be correct, useful, and even superior to what a more competent person would have produced without assistance. Yet the quality of the output ceases to constitute a direct measure of its author's internal competence.

Evidence of this compression already exists. In customer service, generative assistance produced the largest gains among less experienced and lower-performing workers, whereas the initially best-performing workers obtained smaller benefits. The authors interpret the tool as diffusing among novices practices that had previously been concentrated among the most capable workers.

In higher education, Hausman, Rigbi, and Weisburd find that the availability of AI raised the grades of lower-performing students in particular, compressed their distribution, and reduced the informational value of grades for employers. They also distinguish between the acquisition of AI-specific human capital and a possible loss of traditional human capital.

In a randomized experiment with adults of different educational levels, Cruces, Fernández Meijide, Galiani, Gálvez, and Lombardi found that the performance gap between more- and less-educated participants fell by approximately three quarters when they used AI. This assisted equality did not, however, imply that the underlying differences in ability had disappeared.

These results describe a contemporary compression of performance. This article raises an additional question: what happens when the technology continues to improve while, at the same time, subsequent generations of students and professionals are educated through increasing dependence on that same technology?

The proposed answer is that compression can become dynamic and recursive. AI reduces today's observable difference between users with different capacities. Tomorrow, that reduction may alter incentives to acquire knowledge, educational practices, the value of credentials, and the autonomous capacity of new generations themselves. In turn, a population that is less internally trained will become still more dependent on artificial assistance.

The possible result is not merely that less-prepared individuals rise in productive terms. A dual convergence may occur:

1. an upward convergence of outputs, because less competent users receive increasingly comprehensive assistance;
2. a downward convergence of autonomy, because even the strongest members of new cohorts may practise and develop some internal capacities less intensively.

This theory does not assume that all individuals eventually possess the same knowledge. It makes a different claim: people with very different knowledge may end up producing sufficiently similar outputs for differences in knowledge to cease being visible, scarce, or economically relevant.

> *AI does not need to eliminate expertise in order to devalue it. It may be enough to reduce its observable marginal contribution to the result.*

At this point, however, two opposing forces are at work. On the one hand, automation may make it harder for interns and novices to acquire their first experience by eliminating

precisely the entry-level tasks through which they learned the occupation. The resulting lack of practice and internally developed mental models reduces their ability to distinguish a merely plausible answer from a correct one and to detect the errors made by the artificial intelligence itself. On the other hand, progressively better AI may interpret more accurately what the beginner is seeking, guide the beginner toward results once reserved for experts, and act as a tutor by explaining procedures, identifying errors, proposing checks, and accelerating the acquisition of judgment. AI may therefore weaken the traditional pathway through which expertise is formed while simultaneously opening a new pathway of assisted learning. The outcome will depend on whether technology replaces experience or is deliberately used to build it.

2. Conceptual distinctions

2.1. Internal competence

Internal competence is the set of knowledge, mental models, skills, and criteria that a person can mobilize without receiving the substantive part of the solution from an external system.

It does not require working without any tools. A competent professional may consult legislation, manuals, calculators, or databases. The difference is that the professional retains sufficient capacity to:

- understand the problem;
- select the relevant procedure;
- recognize the relevant elements;
- evaluate the coherence of the result;
- explain its foundations;
- detect substantive errors;
- intervene when the ordinary procedure fails.

Internal competence can be approximated through in-person examinations without generative assistance, oral defenses, the resolution of novel problems, transfer exercises, and delayed retention tests.

2.2. Autonomous performance

Autonomous performance is the quality achieved by the individual when the external generative contribution is substantially restricted.

It is not a perfect measure of intelligence or knowledge. It is affected by time pressure, anxiety, memory, and the type of assessment. Nevertheless, it reveals a dimension that assisted output may conceal: what remains when the infrastructure stops providing the solution.

2.3. Assisted performance

Assisted performance is the quality of the output obtained through the combination of human competence, AI capability, and the quality of the interaction between them.

It may greatly exceed autonomous performance. This increase is productively real: a better document remains better even when part of its quality comes from a machine. It cannot, however, be attributed automatically to the user’s human capital.

2.4. Borrowed performance

Borrowed performance refers to the part of productive capacity that an individual can use but cannot autonomously reproduce or justify.

The term does not necessarily carry a pejorative meaning. Societies have always used externalized knowledge. Writing, mathematical tables, libraries, search engines, and software are forms of cognitive offloading. Psychology has studied this phenomenon under the concept of cognitive offloading: reorganizing a task through external resources in order to reduce its internal cognitive demands.

The difference introduced by generative AI is one of degree and possibly of kind. The tool is not limited to storing or retrieving information. It can structure the problem, propose an answer, simulate criticism, revise the text, and integrate knowledge from several disciplines.

2.5. Assisted extended work

This article uses the expression assisted extended work to refer to reports, essays, projects, research, and other outputs that can be produced over an extended period with access to AI.

The expression deliberate work is avoided because it could be confused with the literature on deliberate practice. Extended work may contain genuine human deliberation, but it may also be delegated to a substantial degree.

2.6. Expertise premium

The expertise premium is the difference in quality, productivity, reliability, or remuneration attributable to the experience of one individual relative to a less specialized individual, holding technological access constant.

The theory does not claim that this premium necessarily disappears. It predicts that it will first decline in ordinary tasks clearly located within the AI capability frontier, while it may persist or increase in exceptional, ambiguous, or high-risk tasks.

3. Antecedents and contribution

3.1. Cognitive offloading and automation

Cognitive offloading predates artificial intelligence. Risko and Gilbert show that individuals modify their environment to reduce internal processing demands and that the decision to offload depends, among other factors, on perceived difficulty and metacognitive assessments of one’s own abilities.

Automation has also historically created problems of skill retention. Bainbridge warned that when the system performs ordinary operations and leaves abnormal situations to the human, the operator may lose precisely the practice needed to respond correctly in those exceptional situations.

Generative AI extends this irony to intellectual work. The user may be required to verify a result whose production and foundations the user has ceased to practise.

3.2. Productivity and compression

The initial evidence shows heterogeneous effects. In customer service, Brynjolfsson, Li, and Raymond observed average productivity gains concentrated mainly among novice

workers and those with lower initial ability. More experienced workers obtained small gains and, on some quality measures, even minor declines.

Compression is not universal. The experiment with consultants conducted by Dell'Acqua and co-authors reveals a jagged technological frontier: AI can improve performance when a task falls within its capabilities and worsen it when the user relies on it beyond that frontier.

Accordingly, the relationship between AI and performance inequality depends on the task, the user's level, the quality of the system, and the possibility of verifying the result.

3.3. Education, learning, and signaling

Grades and credentials simultaneously perform educational and signaling functions. In Spence's classic model, education can communicate information about a worker's expected productivity even when part of its value does not arise directly from a productive improvement.

AI may weaken this informational function. If a grade increasingly reflects the capability of the system used rather than the student's autonomous competence, the employer learns less about the candidate by observing the academic record.

Hausman, Rigbi, and Weisburd find precisely an increase in grades concentrated at the lower end of the distribution, grade compression, and erosion of their signaling value.

The relationship between assistance and learning, however, is not unidirectional. In programming experiments, Shen and Tamkin find that intensive delegation can impair conceptual understanding, code reading, and debugging ability, although some cognitively engaged interaction patterns preserve learning.

Conversely, Lira and co-authors show that AI can improve learning when it provides high-quality personalized examples, demonstrating that assistance does not inevitably imply deskilling.

> *The relevant question, therefore, is not whether AI is used, but which cognitive operations it replaces and which it requires the learner to practise.*

3.4. The capability-comprehension gap

Lin and co-authors use the term Capability–Comprehension Gap for the divergence in which assisted performance improves while the internal models required to explain, verify, or intervene deteriorate.

Siu and Fok observe that experts wish to delegate repetitive search tasks but prefer to retain complex synthesis and interpretation. They also emphasize the tension between reducing cognitive load and preserving the practice required to develop expertise.

These concepts are closely related to the decoupling examined here. The specific contribution of the present model is not to discover that operational capability and comprehension can diverge, but to incorporate that divergence into a temporal trajectory connecting:

- the improvement of AI;
- compression between users;
- the formation of new cohorts;
- permeability between specializations;
- the loss of the informational value of outputs;
- and the economic evolution of the expertise premium.

3.5. Specialists and generalists

Recent work has proposed that AI may favor the generalist by reducing the costs of acquiring and coordinating specialized capabilities on demand.

Other models suggest conditional compression: AI reduces differences among users who already exceed a minimum knowledge threshold, but may provide smaller benefits to those who lack the ability required to evaluate the result.

This observation is central. Convergence need not occur linearly. At an early stage, the expert may use AI far more effectively. Only when the system automates part of formulation, contextualization, and verification does the minimum threshold fall and convergence extend toward less competent users.

4. Basic model

Let i denote an individual, j a task, and t a technological moment.

We define:

- $C_{i,t}$: the individual's internal competence;
- A_t: the general capability of the available AI;
- $H_{i,t}$: human-AI interaction competence;
- K_j: task complexity and specificity;
- V_j: task verifiability;
- R_j: the cost or risk of an error;
- $U_{i,j,t}$: autonomous performance;
- $Y_{i,j,t}$: assisted performance.

Autonomous performance can be expressed as:

$$U_{i,j,t} = g(C_{i,t}, K_j) + \varepsilon_{i,j,t}$$

Assisted performance is:

$$Y_{i,j,t} = f(C_{i,t}, A_t, H_{i,t}, K_j, V_j) + \eta_{i,j,t}$$

In the early stages of adoption, internal competence and interaction competence carry substantial weight. The expert obtains better results because the expert knows what to ask, which information to provide, which answer to reject, and which elements to verify.

As AI improves, it can assume a growing share of these operations. For tasks covered by its capability frontier:

$$\partial Y / \partial A > 0$$

and, as a dynamic hypothesis:

$$\partial^2 Y / (\partial A \, \partial C) < 0$$

The second expression indicates that, as AI capability increases, the marginal dependence of the output on human internal competence declines.

4.1. Performance gap

We define the gap between a high-competence user h and a low-competence user l:

$$\Delta^A{}_{j,t} = Y_{h,j,t} - Y_{l,j,t}$$

The compression hypothesis holds that, for ordinary tasks progressively included within the technological frontier:

$$\partial \Delta^{A}_{j,t} / \partial A_t < 0$$

Both users may improve. Compression does not require the better user to worsen in absolute terms. It is enough for the initially weaker user to improve faster or to reach a functional quality ceiling sooner.

4.2. Gap between output and competence

For each individual:

$$B_{i,j,t} = Y_{i,j,t} - U_{i,j,t}$$

B represents the portion of observable performance that disappears or declines when assistance is removed.

A large gap does not imply that the output is useless. It indicates that the output is a poor measure of the individual's autonomous competence.

4.3. Dynamic formation of competence

Future competence depends on learning and cognitive substitution:

$$C_{i,t+1} = C_{i,t} + L(E_{i,t}, A_t) - D(S_{i,t}, A_t)$$

where:

- $E_{i,t}$ represents cognitively engaged interaction;
- $S_{i,t}$ represents substitutive delegation;
- L is facilitated learning;
- D is depreciation or unrealized development.

AI can increase L when it acts as a tutor, provides feedback, or enables better-adapted practice. It can increase D when it replaces operations that the student needed to practise in order to form internal mental models.

The effect on internal competence is therefore conditional rather than technologically inevitable.

4.4. Signaling value

Let:

$$S_t = \mathrm{Corr}(Y_{i,t}, C_{i,t})$$

When assisted output depends increasingly on AI and less on the individual's competence:

$$\partial S_t / \partial A_t < 0$$

Observable quality may increase while its ability to provide information about who produced it declines.

4.5. Graphical synthesis: dual convergence of output and competence

Figure 1 visually summarizes the core of the model for four educational levels: expert, university graduate, high-school graduate, and basic or compulsory secondary education. Each group retains its own color. The solid line represents internal or autonomous competence; the dotted line represents the final output attainable with technological assistance. By construction, assisted output remains above autonomous competence within each group.

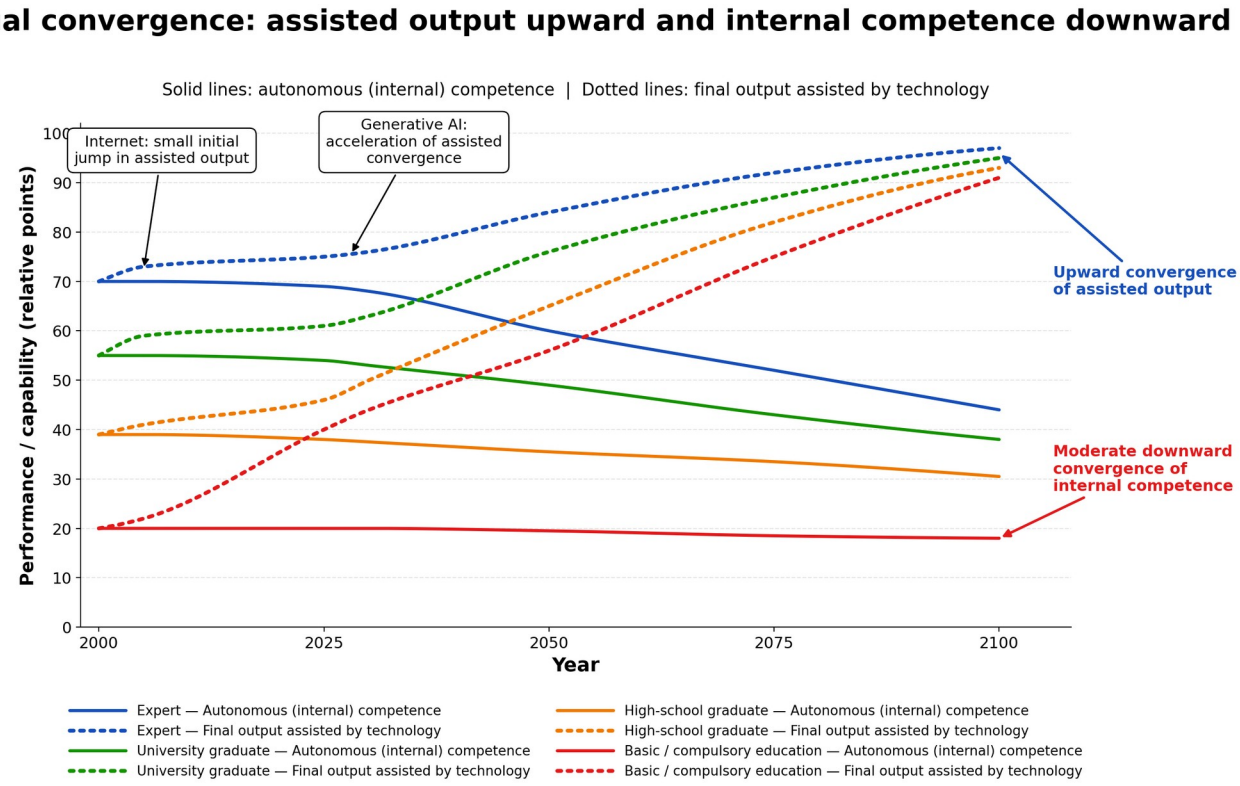
Dual convergence: assisted output upward and internal competence downward
Solid lines: autonomous (internal) competence | Dotted lines: final output assisted by technology
Internet: small initial jump in assisted output
Generative AI: acceleration of assisted convergence
Upward convergence of assisted output
Moderate downward convergence of internal competence
Performance / capability (relative points)
Year
Expert — Autonomous (internal) competence
Expert — Final output assisted by technology
University graduate — Autonomous (internal) competence
University graduate — Final output assisted by technology
High-school graduate — Autonomous (internal) competence
High-school graduate — Final output assisted by technology
Basic / compulsory education — Autonomous (internal) competence
Basic / compulsory education — Final output assisted by technology

Figure 1. Dual convergence: assisted output upward and internal competence downward. Author's own elaboration.

Around the year 2000, the internet produces a small jump in the dotted curves by expanding access to information and tools. It does not, however, produce the same jump in the solid lines: the user must still search, select, understand, and elaborate a large part of the solution. From 2023–2025 onward, generative AI accelerates the rise in assisted output. All groups improve, but the initially lower levels improve proportionally more, so the dotted curves gradually converge toward a high level.

The solid lines follow a different path. After a period of relative stability, they tend to decline moderately when assistance replaces the practice required to acquire and retain skills. Downward convergence is less intense than upward assisted convergence: the expert continues to possess greater autonomous competence, but the distance narrows.

The graph therefore expresses a dual movement: improvement and compression of assisted final output, together with possible erosion and moderate compression of internal competence. The result is a growing decoupling between what a person can deliver with technology and what the person can reproduce, explain, or verify independently. Visible convergence in output is not equivalent to equality of knowledge; it shows precisely that output loses its ability to reveal underlying competence.

> *Society may observe outputs that are increasingly better and more similar, while the internal capacities supporting them become less visible, relatively less valuable, and—under substitutive uses—potentially less robust.*

5. The four phases of dynamic compression

5.1. Phase I: expert complementarity

AI has relevant capabilities, but it is unreliable and requires direction.

The expert retains a substantial advantage because the expert:

- represents the problem more accurately;
- knows the terminology;
- provides appropriate context;
- decomposes the task;
- detects omissions;
- recognizes implausible results;
- distinguishes a plausible answer from a correct answer.

In this phase:

$$\text{Expert} + \text{AI} \gg \text{Non-specialist} + \text{AI}$$

The new technology may even temporarily increase the expertise premium because the expert can extract performance that other users cannot attain.

5.2. Phase II: raising the floor

AI improves the execution of relatively standardized tasks and raises the performance of those who began from a lower baseline in particular.

The best users also improve, but they may be closer to the task's quality ceiling. The absolute or relative difference declines:

$$\Delta^{A}_{t} \downarrow$$

This is the phase best documented by contemporary evidence.

Raising the floor does not imply cognitive equality. Strong differences may remain in examinations without assistance, oral defenses, or the ability to transfer knowledge to novel problems, despite the similarity of submitted work.

5.3. Phase III: automation of supervision

In the third phase, AI ceases to be merely an executor and assumes functions that initially required expertise:

- it requests missing data;
- reformulates deficient instructions;
- checks internal consistency;
- consults sources;
- compares alternatives;
- flags uncertainty;
- performs critical review;
- adapts the output to professional standards;
- proposes verification procedures;
- iteratively corrects its own answers.

Human expertise does not lose relevance merely because the machine executes more, but because it begins to automate what distinguished the expert as a user of the machine.

In this phase:

$$\text{Expert} + \text{AI} \approx \text{Competent generalist} + \text{AI}$$

for a growing set of ordinary tasks.

The competent generalist may become the specialist's closest competitor. The generalist possesses reasoning ability, broad culture, verbal capacity, and transferable judgment, while AI supplies part of the domain-specific knowledge at low cost.

5.4. Phase IV: institutional decoupling

When output quality no longer reveals internal competence sufficiently:

- grades lose discriminatory power;
- academic work loses evidentiary value;
- credentials must be reinterpreted;
- firms face greater difficulty inferring capabilities;
- the wage premium for routinely reproducible knowledge declines;
- the importance of supervised tests, reputation, responsibility, and experience demonstrated in real situations increases.

At this stage, two different but connected mechanisms must be distinguished. The first is the reduced necessity or lower relative value of learning: if AI enables less competent users to produce outputs close to those of experts, additional knowledge yields a smaller marginal advantage and loses part of its productive, wage, and signaling value. The second is the reduced incentive to learn: when the remuneration and recognition attached to educational effort decline, so does the willingness to study intensively, practise for years, and bear the costs of acquiring deep expertise. Reduced necessity is a technological and productive effect; reduced incentive is an economic and institutional effect.

The two mechanisms may reinforce one another recursively: a lower expertise premium reduces investment in study and experience; that lower investment produces fewer deeply competent experts; the shortage of experts increases dependence on AI; and greater dependence further reduces the observable difference between expert and non-expert. Compression of performance may thus also become compression of the incentives to develop capacities that, paradoxically, remain necessary to audit and correct the system.

> *The institution can no longer assume that an excellent output necessarily demonstrates that its author understands what is being presented.*

6. The cohort effect

The theory’s principal contribution is to incorporate generational change explicitly.

Experts currently active acquired much of their competence before having access to AI capable of producing complete outputs. They had to practise operations that new cohorts can delegate from the outset:

- search and selection;
- argument structuring;
- calculation;
- writing;
- programming;
- preliminary diagnosis;
- memory retrieval;
- comparison of hypotheses;
- detection of inconsistencies.

This may create an asymmetry between cohorts.

A graduate from a new generation may produce better assisted work than a graduate from an earlier generation and, at the same time, show lower autonomous performance on certain internal assessments.

The relevant comparison is not only:

Excellent student > Mediocre student

It is also:

whether the strongest graduate in a future cohort maintains, exceeds, or falls below the autonomous competence of the strongest graduate in an earlier cohort,

when both are assessed without the infrastructure characteristic of their respective educational periods.

The strong hypothesis is not that new graduates necessarily know less in absolute terms. They may develop new competences, learn more rapidly, and gain access to an intellectual breadth that was previously impossible. The more limited hypothesis is:

> *When AI systematically replaces a cognitive operation during education, the autonomous competence associated with that operation will tend to develop less unless the educational design introduces compensatory practices.*

This possibility should not be confused with a necessary conclusion. New generations may develop more competences, learn faster, and gain access to an intellectual breadth that was previously impossible; it is also possible that these gains will not compensate for the skills they cease to practise. The model must also distinguish absolute loss from relative devaluation. A new graduate's autonomous competence may decline, but it may also remain stable or even increase and still lose relative value because less competent users, assisted by AI, approach the graduate's results. In that case, the problem would not be that the expert knows less, but that the expert's marginal effort produces a smaller difference relative to those who did not make the same educational investment.

Compression may also erode the academic system's capacity to signal and reward that additional effort. The closer observable outputs become, the harder it is for a grade, paper, or credential to distinguish between someone with deep internal competence and someone who reaches a similar result through assistance. The relative depreciation of expertise may therefore occur even without absolute deskilling.

The cohort effect therefore depends on the pedagogical regime.

Substitutive regime

AI → less practice → lower competence → greater dependence

Tutorial regime

AI → better feedback → more effective practice → greater competence → more critical use

The same technology can activate either loop.

7. Testable hypotheses

H1. Average improvement in assisted performance

The average quality of assisted outputs increases as AI capability improves.

H2. Within-group compression

In standardizable tasks, the assisted-performance difference between individuals with high and low initial competence declines as AI capability increases.

H3. Persistence of the autonomous gap

Compression in assisted output will be greater than compression observed in assessments without assistance.

H4. Output-competence decoupling

The correlation between the quality of submitted output and autonomous competence will decline as the system's generative contribution increases.

H5. Nonlinear trajectory of the expertise premium

The expert advantage may initially increase when AI requires direction and subsequently decline when the system also automates formulation and verification.

The predicted trajectory may take an inverted-U shape:

Expertise premium: ↑ → maximum → ↓

H6. Specialist-generalist convergence

Generalists with high cognitive ability and interaction competence will approach specialists' performance more rapidly than complete novices.

H7. Reduction of the access threshold

The minimum knowledge required to produce an acceptable output will decline as AI incorporates better contextualization, self-correction, and verification mechanisms.

H8. Conditional cohort effect

Cohorts educated under regimes of substitutive delegation will display lower autonomous competence in delegated operations than earlier cohorts or groups educated through tutorial interaction.

H9. Loss of signaling value

Grades, assignments, and credentials based mainly on unsupervised production will lose their ability to predict autonomous performance, subsequent learning, and unaided action.

H10. Shift in the expertise premium

The economic value of the expert will shift from ordinary production toward goal definition, novel problems, uncertainty assessment, responsibility, exceptional intervention, consequence management, and the detection of rare and serious errors.

The hypothesis of a shifting premium should also extend beyond cognitive expertise itself. As symbolic production, analysis, writing, programming, and other white-collar tasks become cheaper, the relative value of executive, manual, bodily, material, and on-site capacities that are not automated at the same pace may rise. This execution premium or blue-collar premium will not benefit all manual work equally: it will depend on scarcity, difficulty of robotization, complementarity with AI, responsibility, demand, and bargaining power.

H11. Temporal tendency of autonomous competence

When AI persistently replaces the cognitive operations required to acquire and maintain a skill, average internal or autonomous competence will tend to stagnate or decline over time. The effect will be stronger in cohorts and settings characterized by substitutive use, and weaker or reversed when AI acts as a tutor, increases cognitively engaged practice, and reinforces retention and transfer.

8. From productive compression to social devaluation

The social value of expert knowledge can be divided into three components.

8.1. Productive value

This is the real improvement that the expert contributes to the result.

As long as:

$$Y_{\text{expert}} - Y_{\text{non-expert}} > 0$$

there is a differential productive contribution.

AI reduces that contribution when it enables the non-expert to produce a sufficiently similar result.

8.2. Scarcity value

A capability receives high remuneration when few people can provide it.

If access to AI enables many people to produce functionally equivalent results, the scarcity of the output declines even if the expert's internal knowledge remains rare.

The market mainly remunerates the scarcity of useful results, not the subjective difficulty of acquiring a capability that no longer appreciably changes those results.

8.3. Signaling value

An output has signaling value when it allows characteristics of its producer to be inferred.

If almost any student can submit an excellent essay, the essay's excellence conveys less information about the student. If any professional can generate a technically competent first draft, that draft conveys less information about the professional's knowledge.

> *AI may devalue knowledge before eliminating it because it can destroy its productive scarcity and social visibility.*

8.4. Additional work and its social return

Suppose that the expert devotes additional years to acquiring knowledge and additional hours to perfecting an output. If that investment only marginally improves a result that another user obtains through AI, the additional work may retain intellectual, personal, or safety value while losing economic recognition.

Society does not directly observe effort or all accumulated knowledge. It observes outputs, decisions, and consequences.

When expert effort ceases to generate a perceptible difference:

$$\text{Marginal social value of effort} \downarrow$$

This does not mean that the effort is useless. It means that it is no longer easily appropriable, demonstrable, or remunerable.

9. Why expertise does not disappear completely

The theory does not predict universal equalization.

9.1. Jagged technological frontier

AI may master some tasks and fail at other apparently adjacent ones. Compression will be greater within the technological frontier and smaller, zero, or even reversed outside it.

9.2. Risk tails

Two users may achieve similar average quality while differing radically in the probability of committing a catastrophic error.

The expert may retain an advantage that is limited in frequency but enormous in consequences.

Equality of means does not imply equality in:

- variance;
- robustness;
- extreme errors;
- adaptation to novel circumstances;
- recognition of uncertainty.

9.3. Responsibility

Even when AI produces the solution, institutions need to assign responsibility.

A physician, lawyer, engineer, or auditor does more than generate answers. The professional assumes consequences, interprets rules, answers to third parties, and must justify decisions.

Responsibility may become a source of value separate from cognitive monopoly.

9.4. Selection of objectives

AI may optimize means without legitimately determining ends.

Choosing which problem to solve, which risk to accept, which interest to prioritize, and which outcome to regard as desirable continues to involve normative, institutional, and political judgment.

9.5. Genuinely novel production

The expert advantage may persist in problems whose representation is not yet available in previous examples or in which quality depends on tacit knowledge, situated experience, and the construction of new concepts.

The appropriate prediction is therefore not the immediate disappearance of the specialist, but a progressive reduction in the space of tasks in which specialist knowledge produces an observable advantage.

10. Educational implications

10.1. Submitted work is no longer sufficient

Universities cannot infer internal competence exclusively from work produced outside supervised conditions.

Nor should the response consist of universally prohibiting AI. The ability to work with it will be a genuine professional competence.

The following should be assessed separately:

1. what the student can produce with AI;
2. what the student understands;
3. what the student can do without it;
4. how the student verifies results;
5. how the student responds when AI fails.

10.2. Dual assessment

A reasonable architecture would combine:

- professionally rigorous assisted work;
- an internal examination without assistance;
- oral defense;
- process audit;
- transfer exercise;
- error-detection test;
- delayed retention assessment.

The objective is not to force students always to work without technology, but to prevent the institution from confusing augmented capability with possessed competence.

Dual assessment must explicitly separate two outcomes and track their temporal evolution independently. The first is autonomous or internal competence, measured through tests in which generative assistance is restricted. The second is assisted performance or final-output quality, measured while transparently permitting AI and other tools. Both dimensions are legitimate and relevant, but they answer different questions: one indicates what the person understands and can execute independently; the other indicates what the person can produce within the real technological infrastructure.

This separation reveals trajectories that an aggregate grade would conceal: assisted output may improve while internal competence remains stable; both dimensions may improve; or output may increase while autonomous capability declines. Universities and experiments should report both metrics rather than prematurely converting them into a single result.

The sporting analogy is useful: it would be like allowing an athlete first to compete in a race without doping and then in another race in which certain forms of doping were expressly permitted by the rules. Both tests may be legitimate if they are clearly identified and assessed separately; what would be misleading is to present the result of the race with permitted doping as though it demonstrated exactly the same autonomous capability as the race without doping.

10.3. Multidimensional certification

Future credentials should distinguish between:

- autonomous competence;
- AI-interaction competence;
- verification capacity;
- augmented professional performance;
- responsibility and judgment.

A single academic average may conceal radically different combinations.

Diplomas, transcripts, and certifications should separately display grades obtained in autonomous-knowledge tests and those obtained in assisted work or final products.

Employers could then interpret each result as a signal of a specific capability, without confusing internal mastery, ability to operate with AI, and quality of the delivered output.

Transparent certification should make it possible to know simultaneously how much a candidate can produce with the available tools and what knowledge the candidate retains, understands, and can mobilize without them. The purpose is not necessarily to privilege one modality over the other, but to prevent distinct capabilities from being hidden within a single figure and to preserve the informational value of credentials.

10.4. Protecting the training pathway

If AI automates the entry-level tasks through which novices acquired experience, organizations may find themselves without mechanisms for producing future experts.

The problem is not merely to protect jobs. It is to preserve a pathway of learning.

Entry-level tasks may appear inefficient from the perspective of immediate production while being indispensable from the perspective of reproducing human capital.

11. Labor-market implications

11.1. Lower value of ordinary specialization

Specialization based on retrieving, combining, and presenting relatively stable information will be more vulnerable than specialization based on responsibility, exceptional experience, or knowledge that is difficult to codify.

11.2. Relative rise of the generalist

The generalist can use AI as an infrastructure for specialization on demand.

The generalist will not necessarily replace the highest-level expert, but may replace several intermediate specialists in sufficiently standardized tasks.

11.3. Polarization within expertise

The category of “expert” may divide into:

- experts whose ordinary output becomes commoditized;
- frontier experts who define problems, validate systems, or intervene in exceptional cases.

The average specialization premium may fall while the remuneration of a minority responsible for extreme cases rises.

11.4. New signaling mechanisms

When outputs are easily reproducible, markets will seek alternative signals:

- in-person tests;
- verifiable track record;
- reputation;
- real-world results;
- supervised certifications;
- legal responsibility;
- ability to explain decisions;
- performance under uncertainty.

The loss of one signal does not eliminate the need for selection. It triggers the creation of new signals, potentially at greater cost.

12. Empirical strategy

12.1. Within-cohort university experiment

Students would complete four assessments on the same subject:

1. an examination without AI;
2. extended work produced with AI;
3. an oral defense of the work;
4. a novel problem without AI several weeks later.

The main variables would be:

- quality;
- time;
- accuracy;
- depth;
- error rate;
- explanatory ability;
- retention;
- transfer.

Prediction:

- dispersion will be lower in assisted work;
- the correlation between assisted work and the internal examination will be lower than the historical correlation between assignments and examinations;
- part of the difference will reappear in the defense, the transfer task, and the delayed test.

12.2. Longitudinal cohort comparison

Cohorts educated before and after the intensive adoption of AI would be compared using equivalent tests of autonomous competence.

The analysis must control for:

- student selection;
- curricular changes;
- test difficulty;
- demographic composition;
- teaching quality;
- institutional disruptions.

The deskilling prediction would be confirmed only if the decline appeared specifically in substituted operations rather than across all capacities.

12.3. Specialist-generalist experiment

Participants would include:

- domain experts;
- high-ability generalists;
- users without relevant experience.

All would work with different generations of AI or with increasing levels of automation.

The theory predicts:

1. a clear initial expert advantage;

2. rapid approximation by the generalist;
3. later convergence by the less competent user when the system automates supervision;
4. persistence of the expert advantage in atypical cases and high-cost errors.

12.4. Signaling experiment

Employers or evaluators would receive identical outputs accompanied by different information concerning:

- grade;
- AI use;
- result of a supervised test;
- oral defense;
- certification of internal competence.

The experiment would measure how valuation changes when evaluators know that the relationship between output and autonomous capability is weak.

12.5. Measuring the expertise premium

The premium should not be measured solely through average quality.

It should include:

$$\text{Premium} = \Delta\,\text{quality} + \Delta\,\text{reliability} + \Delta\,\text{risk management} + \Delta\,\text{responsibility} - \Delta\,\text{cost}$$

The average quality difference may tend toward zero while the reliability difference in extreme cases continues to justify a high premium.

13. Conditions of falsification

The theory would be weakened if evidence systematically showed that:

1. initially stronger users expand their advantage with each new generation of AI;
2. improved assisted performance translates consistently into an equivalent improvement in autonomous performance;
3. AI-intensive cohorts develop equal or greater internal competence in delegated operations;
4. generalists do not approach specialists even in ordinary tasks;
5. automation of verification does not reduce the minimum knowledge threshold;
6. employers continue to extract the same information from assignments and grades;
7. remuneration for expertise rises in all tasks, including fully automatable ones;
8. compression disappears when more advanced tools are used.

The model must also accommodate mixed results. AI may compress some differences and widen others. It may reduce the importance of traditional knowledge and increase the importance of interaction competence. It may raise the productivity of less-prepared users while simultaneously creating a new inequality based on the ability to direct complex systems.

14. Discussion

The optimistic interpretation holds that AI democratizes previously scarce capabilities. From this perspective, the ability of a person to produce an excellent result without years of education is a social gain rather than a loss.

This interpretation is partly correct. Compression can reduce barriers, raise productivity, enable more flexible careers, and extend high-quality services to previously excluded populations.

But democratization of output does not automatically solve three problems.

1. it may conceal who truly understands the result;
2. it may weaken the processes through which people capable of intervening when the system fails are trained;
3. it may erode private incentives to acquire knowledge whose social benefit remains, but whose individual contribution ceases to be visible and remunerable.

There is therefore a possible divergence between individual rationality and collective need.

For the individual, it may be rational not to invest years in a competence that AI reproduces at low cost. Society may nevertheless still need a sufficient number of people capable of understanding, auditing, and reconstructing the procedure.

This problem recalls the classic irony of automation: the better the system performs in ordinary situations, the less practice humans acquire for the exceptional situations in which their intervention becomes essential.

Dynamic compression also requires reconsidering the meaning of “knowing.” A person may not internally possess all the required information and yet be able to mobilize an infrastructure that makes the person productively competent.

It would be incorrect to claim that such a person knows nothing. But it would also be incorrect to attribute individually to that person everything incorporated into the output.

The relevant productive unit shifts from the isolated individual to the assemblage:

$$\text{Human} + \text{AI} + \text{data} + \text{institution}$$

The economic and educational problem is then to determine which part of capability must remain within the individual in order to preserve autonomy, responsibility, learning, and the capacity to intervene.

15. Conclusion

Generative artificial intelligence does not first eliminate expert knowledge. It first eliminates the need for that knowledge to be fully incorporated into every person who produces a competent result.

In a first phase, AI complements the expert and may increase the expert’s advantage. It then raises the performance of less competent users and compresses the distribution of outputs. Later, by automating formulation, contextualization, and verification, it erodes the functions that protected the expertise premium. At the same time, new cohorts may develop some internal capacities less fully when assistance replaces the practice needed to acquire them.

The result is a progressive separation among three elements that were previously more closely associated:

$$\text{knowledge}, \quad \text{output}, \quad \text{social value}$$

Outputs may improve while their ability to reveal knowledge declines. Individuals may become more productive while becoming less autonomous. Experts may continue to know more while that knowledge produces a smaller difference in ordinary production.

The consequence will not necessarily be complete equality. Expertise may retain decisive value in novel problems, extreme cases, uncertainty, responsibility, and catastrophic failures. But its economic domain will narrow as a growing share of tasks moves within the technological frontier.

The central thesis can be summarized as follows:

> *AI produces a dynamic compression of expertise when technological improvement reduces the dependence of output on human competence while, at the same time, substitutive use reduces the autonomous formation of new cohorts. This interaction may cause the outputs of students, generalists, and experts to converge before their internal capacities converge, subsequently weakening the informational, economic, and social value of specialized knowledge.*

The decisive question is no longer merely how much more a person can produce with AI. It is how much a society must continue to understand for itself if it intends to preserve the capacity to judge, correct, and govern what its machines produce.

References

Bainbridge, L. (1983). “Ironies of Automation”. Automatica, 19(6), 775–779.

Brynjolfsson, E., Li, D., and Raymond, L. R. (2023/2026). “Generative AI at Work”. Study of generative assistance and productivity in customer service.

Cruces, G., Fernández Meijide, D., Galiani, S., Gálvez, R. H., and Lombardi, M. (2026). “Does Generative AI Narrow Education-Based Productivity Gaps? Evidence from a Randomized Experiment”. NBER Working Paper 34851.

Dell’Acqua, F. et al. (2026). “Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality”. Organization Science.

Hausman, N., Rigbi, O., and Weisburd, S. (2025). “Generative AI’s Impact on Student Achievement and Implications for Worker Productivity”. Working paper.

Lin, F. et al. (2026). “Human-Centric AI Requires a Minimum Viable Level of Human Understanding”. Preprint.

Lira, B., Rogers, T., Goldstein, D. G., Ungar, L., and Duckworth, A. L. (2025). “Learning Not Cheating: AI Assistance Can Enhance Rather Than Hinder Skill Development”. Preprint.

Risko, E. F., and Gilbert, S. J. (2016). “Cognitive Offloading”. Trends in Cognitive Sciences, 20(9), 676–688.

Shen, J. H., and Tamkin, A. (2026). “How AI Impacts Skill Formation”. Preprint.

Siu, A., and Fok, R. (2025). “Augmenting Expert Cognition in the Age of Generative AI: Insights from Document-Centric Knowledge Work”. Preprint.

Spence, M. (1973). “Job Market Signaling”. The Quarterly Journal of Economics, 87(3), 355–374.